

# Restricted Collapsed Draw: Accurate Sampling for Hierarchical Chinese Restaurant Process Hidden Markov Models

## Abstract

We propose a restricted collapsed draw (RCD) sampler, a general Markov chain Monte Carlo sampler of simultaneous draws from a hierarchical Chinese restaurant process (HCRP) with restriction. Models that require simultaneous draws from a hierarchical Dirichlet process with restriction, such as infinite Hidden markov models (iHMM), were difficult to enjoy benefits of the HCRP due to combinatorial explosion in calculating distributions of coupled draws. By constructing a proposal of seating arrangements (partitioning) and stochastically accepts the proposal by the Metropolis-Hastings algorithm, the RCD sampler makes accurate sampling for complex combination of draws while retaining efficiency of HCRP representation. Based on the RCD sampler, we developed a series of sophisticated sampling algorithms for iHMMs, including blocked Gibbs sampling, beam sampling, and split-merge sampling, that outperformed conventional iHMM samplers in experiments.

## 1 Introduction

Existing sampling algorithms for infinite hidden Markov models (iHMMs, also known as the hierarchical Dirichlet process HMMs) [??] do not use a hierarchical Chinese restaurant process (HCRP) [?], which is a way of representing the predictive distribution of a hierarchical Dirichlet process (HDP) by collapsing, i.e. integrating out, the underlying distributions of the Dirichlet process (DP). While an HCRP representation provides efficient sampling for many other models based on an HDP [??] through reducing the dimension of sampling space, it has been considered rather “awkward” [?] to use an HCRP for iHMMs, due to the difficulty in handling coupling between random variables. In the simplest case, consider step-wise Gibbs sampling from an iHMM defined as  $\pi_k \sim \text{DP}(\beta, \alpha_0)$  and




Figure 1: Step-wise Gibbs sampling in iHMM. Since the Dirichlet process prior is posed on *transitions* in iHMM, resampling  $x_i$  involves taking two transitions,  $x_{i-1} \rightarrow x_i$  and  $x_i \rightarrow x_{i+1}$ , simultaneously. In this case, we consider distribution of two draws  $(x'_i, x'_{i+1})$  with restriction that the draws are consistent with remaining sequence, i.e.,  $x'_{i+1} = x_{i+1}$ .

$\beta \sim \text{GEM}(\gamma)$ . Given  $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_T$ , resampling hidden state  $x_i$  actually consists of two draws (Figure 1),  $x'_i \sim \pi_{x_{i-1}}$  and  $x'_{i+1} \sim \pi_{x'_i}$ , under the restriction  $(x'_i, x'_{i+1}) \in C$  that these draws are consistent with the following sequence, i.e.,  $C = \{(x'_i, x'_{i+1}) | x'_{i+1} = x_{i+1}\}$ . Under the HCRP, the two draws are coupled even if  $x_{i-1} \neq x'_i$ , because distributions  $\pi_{x_{i-1}}, \pi_{x'_i}$  as well as the base measure  $\beta$  are integrated out in an HCRP, and coupling complicates sampling from the restricted distribution.

To generalize, the main part of the difficulty is to obtain a sample from a restricted joint distribution of simultaneous draws from collapsed distributions, which we call *restricted collapsed draw* (RCD). Consider resampling  $L$  draws simultaneously,  $\mathbf{x} = (x_{j_1 i_1}, \dots, x_{j_L i_L})$ , from the respective restaurants  $\mathbf{j} = (j_1, \dots, j_L)$ , when we have a restriction  $C$  such that  $\mathbf{x} \in C$ . Step-wise Gibbs sampling from iHMM can be fitted into RCD with  $L = 2$  by allowing

restaurant index  $j_2$  to be dependent on the preceding draw  $x_{j_1 i_1}$ .

In this paper, we point out that it is not enough to consider the distribution of draws. Since the HCRP introduces an additional set of latent variables  $s$  that accounts for the seating arrangements of the restaurants, we have to compute an exact distribution of  $s$  as well, under the restriction. We want to perform sampling from the following conditional distribution,

$$p(\mathbf{x}, \mathbf{s} | C) = \frac{1}{Z_C} \mathbb{I}[\mathbf{x} \in C] p(\mathbf{x}, \mathbf{s}) \quad , \quad (1)$$

where  $Z_C$  is a normalization constant and  $\mathbb{I}$  is the indicator function, whose value is 1 if the condition is true and 0 otherwise. Although non-restricted probability  $p(\mathbf{x}, \mathbf{s})$  can be easily calculated for a given  $\mathbf{x}$  and  $\mathbf{s}$ , calculating the normalization constant  $Z_C$  leads to a combinatorial explosion in terms of  $L$ .

To solve this issue, we propose the restricted collapsed draw (RCD) sampler, which provides accurate distributions of simultaneous draws and seating arrangements from HCRP. The RCD sampler constructs a proposal of seating arrangements using a given proposal of draws, and the pair of proposals are stochastically accepted by the Metropolis-Hastings algorithm [?]. Since the RCD sampler can handle any combination of restricted collapsed draws simultaneously, we were able to develop a series of sampling method for HCRP-HMM, including a blocked collapsed Gibbs sampler, a collapsed beam sampler, and a split-merge sampler for HCRP-HMM. Through experiments we found that our collapsed samplers outperformed their non-collapsed counterparts.

## 2 HCRP representation for iHMM

### 2.1 Infinite HMM

An infinite hidden Markov model (iHMM) [??] is defined over the following HDP:

$$G_0 \sim \text{DP}(\gamma, H) \quad G_k \sim \text{DP}(\alpha_0, G_0) \quad , \quad (2)$$

To see the relation of this HDP to the transition matrix  $\pi$ , consider the explicit representation of parameters:

$$G_0 = \sum_{k'=1}^{\infty} \beta_{k'} \phi_{k'} \quad G_k = \sum_{k'=1}^{\infty} \pi_{k'k} \phi_{k'} \quad , \quad (3)$$

where transition probability  $\pi_k$  is given as  $\pi_k \sim \text{DP}(\alpha_0, \beta)$ ,  $\beta \sim \text{GEM}(\gamma)$  is the stick-breaking construction of DPs [?], and  $\phi_k \sim H$ .

A formal definition for the HDP based on this representation is:

$$\beta | \gamma \sim \text{GEM}(\gamma) \quad \pi_j | \alpha_0, \beta \sim \text{DP}(\alpha_0, \beta) \quad (4)$$

$$x_{ji} | (\pi_k)_{k=1}^{\infty} \sim \pi_j \quad \phi_k \sim H \quad y_{ji} | x_{ji} \sim F(\phi_{x_{ji}}) \quad , \quad (5)$$

Given an HDP and initial state  $x_0$ , we can construct an infinite HMM by extracting a sequence of draws  $x_i$  as  $x_i = x_{x_{i-1}i}$ , and corresponding observations  $y_i = y_{x_i i}$ . Figure 2 shows a graphical representation of the iHMM.

### 2.2 HCRP-HMM

As another way of representing HDP in iHMM (Eq. 2), we introduce a hierarchical Chinese restaurant process (HCRP, also known as the Chinese restaurant franchise), which does not need to sample the transition distribution  $\pi$  and its base measure  $\beta$  in Eq. (4):

$$k_{jt} | \gamma \sim \text{CRP}(\gamma) \quad t_{ji} | \alpha_0 \sim \text{CRP}(\alpha_0) \quad (6)$$

$$x_{ji} = k_{jt_{ji}} \quad (7)$$

$$\phi_k \sim H \quad y_{ji} | x_{ji}, \phi \sim F(\phi_{x_{ji}}) \quad . \quad (8)$$

Using the Chinese restaurant metaphor, we say that customer  $i$  of restaurant  $j$  sits at table  $t_{ji}$ , which has a dish of an index  $k_{jt_{ji}}$ .

To understand connection between HDP and HCRP, consider a finite model of grouped observations  $x_{ji}$ , in which each group  $j$  choose a subset of  $M$  mixture components from a model-wide set of  $K$  mixture components:

$$\beta | \gamma \sim \text{Dir}(\gamma/K, \dots, \gamma/K) \quad k_{jt} | \beta \sim \beta \quad (9)$$

$$\tau_j | \alpha_0 \sim \text{Dir}(\alpha_0/M, \dots, \alpha_0/M) \quad t_{ji} | \tau_j \sim \tau_j \quad (10)$$

As  $K \rightarrow \infty$  and  $M \rightarrow \infty$ , the limit of this model is HCRP; hence the infinite limit of this model is also HDP. Equation (6) is derived by taking the infinite limit of  $K$  and  $M$  after integrating out  $\beta$  and  $\tau$  in Eqs. (9) and (10). The distribution  $\pi_j$  in Eq. 4 can be derived from  $\tau_j$  and  $k_j$  as follows:

$$\pi_j = \sum_t \tau_{jt} \delta_{k_{jt}} \quad . \quad (11)$$

To consider sampling of  $x_{ji}$  using HCRP (Eqs. 7 and 8), we use count notation  $n_{jtk}$  as the number of customers in restaurant  $j$  at table  $t$  serving the dish of the  $k$ -th entry, and  $m_{jk}$  as the number of tables in restaurants the  $j$  serving the dish of the  $k$ -th entry. We also use dots for marginal counts (e.g.,  $m_{\cdot k} = \sum_j m_{jk}$ ). Then, we sample table index  $t_{ji}$  from the following distribution:

$$p(t_{ji} = t | t_{j1}, \dots, t_{j,i-1}) = \frac{n_{jt\cdot}}{n_{j\cdot\cdot} + \alpha_0} \quad (12)$$

$$p(t_{ji} = t^{new} | t_{j1}, \dots, t_{j,i-1}) = \frac{\alpha_0}{n_{j\cdot\cdot} + \alpha_0} \quad . \quad (13)$$

When  $t_{ji} = t^{new}$  (i.e., the customer sits at a new table), we need to sample  $k_{jt^{new}}$ , whose distribution is:

$$p(k_{jt} = k | k_{11}, \dots) = \frac{m_{\cdot k}}{m_{\cdot\cdot} + \gamma} \quad (14)$$

$$p(k_{jt} = k^{new} | k_{11}, \dots) = \frac{\gamma}{m_{\cdot\cdot} + \gamma} \quad . \quad (15)$$

These variables determine the new sample  $x_{ji} = k_{jt_{ji}}$ . Since  $x_{ji}$  does not uniquely determine the state of the HCRP model, we need to keep latent variables  $t_{ji}$  and




Figure 2: Graphical Representation of iHMM.

$k_{jt}$  for subsequent sampling. We will denote  $s^{(j)} = (t_{j1}, t_{j2}, \dots)$  as the seating arrangement in restaurant  $j$ ,  $s^{(0)} = (k_{11}, k_{12}, \dots, k_{21}, \dots)$  as the seating arrangement in the root restaurant, and  $\mathbf{s}$  as the collection of all seating arrangements, corresponding to the sampled model state. In Bayesian inference based on sampling, we need a procedure to sample the latent variables, given the value of new draw  $x_{ji}$  and the seating arrangements for other draws  $\mathbf{s}$ , which is called as addCustomer.

Construction of HCRP-HMM is the same as iHMM, i.e., extracting a sequence of draws  $x_i$  given  $x_0$  as  $x_i = x_{x_{i-1}i}$ , and corresponding observations  $y_i = y_{x_{i-1}i}$ .

### 3 Restricted Collapsed Draw Sampler

What we want is a sampling algorithm for HCRP-HMM. As described in the Introduction, the problem can be reduced to an algorithm for sampling from  $p(\mathbf{x}, \mathbf{s}|C)$ , i.e., the distribution of restricted collapsed draw with seating arrangements (Eq. 1).

Our idea is to apply the Metropolis-Hastings algorithm [?] to the seating arrangements, which stochastically accepts the proposal distribution of seating arrangements. Although it is hard to directly give proposal distribution  $q(\mathbf{s})$  of seating arrangements, our method constructs  $q(\mathbf{s})$  by combining  $q_{\mathbf{x}}(\mathbf{x})$  with  $q_{\mathbf{s}}(\mathbf{s}|\mathbf{x})$ , another proposal of seating arrangements given the proposed draws, which is based on the addCustomer procedure that is standardly used in Gibbs sampling of HCRP.

#### 3.1 Overall sampling

The Metropolis-Hastings algorithm provides a way of constructing an MCMC sampler using unnormalized probability value  $\tilde{p}(z)$ . After sampling  $z^*$  from proposal distribution  $q(z^*|z)$ , the algorithm computes acceptance probability  $R$ :

$$R = \min \left( 1, \frac{\tilde{p}(z^*)}{\tilde{p}(z^{old})} \frac{q(z^{old}|z^*)}{q(z^*|z^{old})} \right). \quad (16)$$

Then the result  $z^{new} = z^*$  with probability  $R$ , and  $z^{new} = z^{old}$  otherwise. Repeating this process constitutes an MCMC sampler from required distribution  $p(z) \propto \tilde{p}(z)$ ,

Within the context of HCRP, sample space  $\mathbf{z}$  consists of draws  $\mathbf{x}$  and seating arrangement  $\mathbf{s}$ . From Eq. (1), we can use the non-restricted probability of draws  $p(\mathbf{x}, \mathbf{s})$  as unnormalized probability value  $\tilde{p}(z)$ , but it is not easy to provide a proposal for joint distribution  $q(\mathbf{x}^*, \mathbf{s}^*)$ .

Our idea is to factorize the proposal distribution as:

$$q(\mathbf{x}^*, \mathbf{s}^*|\mathbf{s}_0) = q_{\mathbf{x}}(\mathbf{x}^*|\mathbf{s}_0) \cdot q_{\mathbf{s}}(\mathbf{s}^*|\mathbf{x}^*, \mathbf{s}_0). \quad (17)$$

First factor  $q_{\mathbf{x}}$  is the proposal distribution of the draws. Second factor  $q_{\mathbf{s}}$  is the proposal distribution of the seating arrangements given the proposal draws. We use the result of the addCustomer procedure, which stochastically updates the seating arrangements, as the proposal distribution of the seating arrangements.

#### 3.2 Computing Factors

The following describes each factor in  $R$  and its computation.

**True Probability**  $p(\mathbf{x}, \mathbf{s})$  in Eq. (1) is the joint probability of all draws  $x_{ji}$ :

$$p(\mathbf{x}, \mathbf{s}) = \prod_j p(s^{(j)}) \cdot p(s^{(0)}) \quad (18)$$

where

$$p(s^{(j)}) = \frac{\Gamma(\alpha_0)}{\Gamma(\alpha_0 + n_{j..})} \cdot \alpha_0^{m_{j..}} \cdot \prod_t \Gamma(n_{jt..}) \quad (19)$$

$$p(s^{(0)}) = \frac{\Gamma(\gamma)}{\Gamma(\gamma + m_{..})} \cdot \gamma^K \cdot \prod_t \Gamma(m_{..k}) \quad (20)$$

and  $\Gamma$  is the Gamma function. This is the product of the probabilities of seating arrangements (Eqs. 12 to 15) for each customer.

In practice, we only need to calculate probabilities that account for the change in seating from  $\mathbf{s}_0$ , because the prob-

ability for unchanged customers is cancelled out through reducing the fraction in  $R$ . Let  $\mathbf{s}_0$  be the seating arrangement for the unchanged customers, then

$$\frac{p(\mathbf{s}^*)}{p(\mathbf{s}^{old})} = \frac{p(\mathbf{s}^*|\mathbf{s}_0)}{p(\mathbf{s}^{old}|\mathbf{s}_0)} . \quad (21)$$

In fact,  $p(\mathbf{x}^*, \mathbf{s}^*|\mathbf{s}_0)$  is easily calculated along with `addCustomer` operations:

$$p(\mathbf{x}^*, \mathbf{s}^*|\mathbf{s}_0) = p(x_1^*, \mathbf{s}_1^*|\mathbf{s}_0) p(x_2^*, \mathbf{s}_2^*|\mathbf{s}_1^*) \cdots p(x_L^*, \mathbf{s}_L^*|\mathbf{s}_{L-1}^*) . \quad (22)$$

Here,  $p(x_\ell, \mathbf{s}_\ell|\mathbf{s}_{\ell-1})$  is probability  $p(x_{j_\ell i_\ell}, t_{j_\ell i_\ell}|j_\ell, \mathbf{s}_{\ell-1})$  of obtaining seating arrangement  $\mathbf{s}_\ell$  as a result of drawing a sample from restaurant  $j$ :

$$p(x_{ji} = k, t_{ji} = t|j, \mathbf{s}) = \frac{1}{n_{j..} + \alpha_0} \times \begin{cases} n_{jtk} & n_{jtk} \geq 1 \\ \alpha_0 \cdot \frac{m_{..k}}{m_{..} + \gamma} & n_{jtk} = 0, m_{..k} \geq 1 \\ \alpha_0 \cdot \frac{\gamma}{m_{..} + \gamma} & n_{jtk} = 0, m_{..k} = 0 \end{cases} \quad (23)$$

The same applies to the calculation of  $p(\mathbf{s}^{old}|\mathbf{s}_0)$ , which can be done along with `removeCustomer` operations.

**Proposal Distribution of Draws**  $q(\mathbf{x})$  can be anything as long as it is ergodic within restriction  $C$ . To increase the acceptance probability, however, it is preferable for the proposal distribution to be close to the true distribution. We suggest that a good starting point would be to use a joint distribution composed of the predictive distributions of each draw, as has been done in the approximated Gibbs sampler [?]:

$$q_{\mathbf{x}}(\mathbf{x}) = \mathbb{I}[\mathbf{x} \in C] \prod_{i=1}^L p(x_i|\mathbf{s}_0) . \quad (24)$$

We will again discuss the proposal distribution of draws for the HCRP-HMM case in Section 4.

**Proposal Distribution of Seating Arrangements**  $q_{\mathbf{s}}(\mathbf{s}^*|\mathbf{x}, \mathbf{s}_0)$ , is the product of the probabilities for each operation of adding a customer:

$$q_{\mathbf{s}}(\mathbf{s}^*|\mathbf{x}^*, \mathbf{s}_0) = q_{\mathbf{s}}(\mathbf{s}_1^*|x_1^*, \mathbf{s}_0) q_{\mathbf{s}}(\mathbf{s}_2^*|x_2^*, \mathbf{s}_1^*) \cdots q_{\mathbf{s}}(\mathbf{s}_\ell^*|x_\ell^*, \mathbf{s}_{\ell-1}^*) . \quad (25)$$

Here,  $q_{\mathbf{s}}(\mathbf{s}_\ell|\mathbf{x}_\ell, \mathbf{s}_{\ell-1}) = p(t_{j_\ell i_\ell}|x_{j_\ell i_\ell}, j_\ell, \mathbf{s}_{\ell-1})$ , i.e., the probability of obtaining seating arrangement  $\mathbf{s}_\ell$  as a result of the `addCustomer`( $x_{j_\ell i_\ell}, j_\ell, \mathbf{s}_{\ell-1}$ ) operation.

$$p(t_{ji} = t|x_{ji} = k, j, \mathbf{s}) = \begin{cases} \frac{n_{jtk}}{n_{j..k} + \alpha_0 \frac{m_{..k}}{m_{..} + \gamma}} & n_{jtk} \geq 1 \wedge k_{jt} = k \\ \frac{\alpha_0 \frac{m_{..k}}{m_{..} + \gamma}}{n_{j..k} + \alpha_0 \frac{m_{..k}}{m_{..} + \gamma}} & n_{jtk} = 0 \wedge m_{..k} > 0 \\ 1 & n_{jtk} = 0 \wedge m_{..k} = 0 \end{cases} . \quad (26)$$

### 3.3 Simplification

Paying attention to the fact that both Eqs. (23) and (27) are calculated along a series of `addCustomer` calls, we introduce factors

$$r_\ell^* = \frac{p(x_\ell^*, \mathbf{s}_\ell|\mathbf{s}_{\ell-1})}{q_{\mathbf{s}}(\mathbf{s}_\ell|x_\ell^*, \mathbf{s}_{\ell-1})} \quad r_\ell^{old} = \frac{p(x_\ell^{old}, \mathbf{s}_\ell|\mathbf{s}_{\ell-1})}{q_{\mathbf{s}}(\mathbf{s}_\ell|x_\ell^{old}, \mathbf{s}_{\ell-1})} \quad (28)$$

to simplify the calculation of  $R$  as:

$$R = \min \left( 1, \frac{p(\mathbf{s}^*)}{p(\mathbf{s}^{old})} \frac{q_{\mathbf{s}}(\mathbf{s}^{old}|\mathbf{x}^{old}, \mathbf{s}_0)}{q_{\mathbf{s}}(\mathbf{s}^*|\mathbf{x}^*, \mathbf{s}_0)} \frac{q_{\mathbf{x}}(\mathbf{x}^{old})}{q_{\mathbf{x}}(\mathbf{x}^*)} \right) = \min \left( 1, \frac{r(\mathbf{x}^*, \mathbf{s}^*|\mathbf{s}_0)}{r(\mathbf{x}^{old}, \mathbf{s}^{old}|\mathbf{s}_0)} \frac{q(\mathbf{x}^{old})}{q(\mathbf{x}^*)} \right) , \quad (29)$$

where

$$r(\mathbf{x}^*, \mathbf{s}^*|\mathbf{s}_0) = \frac{p(\mathbf{x}^*, \mathbf{s}^*|\mathbf{s}_0)}{q_{\mathbf{s}}(\mathbf{s}^*|\mathbf{x}^*, \mathbf{s}_0)} = \frac{p(x_1^*, \mathbf{s}_1^*|\mathbf{s}_0)}{q_{\mathbf{s}}(\mathbf{s}_1^*|x_1^*, \mathbf{s}_0)} \cdots \frac{p(x_L^*, \mathbf{s}_L^*|\mathbf{s}_{L-1}^*)}{q_{\mathbf{s}}(\mathbf{s}_L^*|x_L^*, \mathbf{s}_{L-1}^*)} = r_1^* \cdot r_2^* \cdots r_L^* . \quad (30)$$

Surprisingly, assigning Eqs. (23) and (27) into Eq. (28) reveals that  $r_\ell^*$  is equal to  $p(x_{j_\ell i_\ell} = x_\ell^*|\mathbf{s}_{\ell-1}^*)$ , i.e., the probability of new customer  $x_{j_\ell i_\ell}$  at restaurant  $j_\ell$  eating dish  $x_\ell^*$ :

$$p(x_{ji} = k|\mathbf{s}) = \frac{n_{j..k} + \alpha_0 \frac{m_{..k}}{m_{..} + \gamma}}{n_{j..} + \alpha_0} \quad (31)$$

$$p(x_{ji} = k^{new}|\mathbf{s}) = \frac{\alpha_0 \frac{m_{..k}}{m_{..} + \gamma}}{n_{j..} + \alpha_0} . \quad (32)$$

In other words, calculation of the accept ratio does not use  $t_{ji}$  (the table index of each customer), despite the fact that the values of  $t_{ji}$  are being proposed;  $t_{ji}$  will indirectly affect the accept ratio by changing subsequent draw probabilities  $p(x_{\ell+1}^*|\mathbf{s}_\ell^*), p(x_{\ell+2}^*|\mathbf{s}_{\ell+1}^*), \dots$  through modifying  $n_{jtk}$  and  $m_{jk}$ , i.e., the number of customers and tables.

It is now clear that, as done in some previous work [?], we can save storage space by using an alternative representation for seating arrangements  $\mathbf{s}$ , in which the table indices of each customer  $t_{ji}$  are forgotten but only the numbers of customers  $n_{jt.}$ ,  $k_{jt}$  and  $m_{jk}$  are retained. The only remaining reference to  $t_{ji}$  in the `removeCustomer` procedure can be safely replaced by sampling.

However, it should be noted that we have to revert to original seating assignment  $\mathbf{s}^{old}$  whenever the proposal is rejected. Putting the old draws  $\mathbf{x}^{old}$  back into  $\mathbf{s}_0$  by using the `addCustomer` procedure again will lead sampling to an incorrect distribution of seating assignments, and consequently, an incorrect distribution of draws.

Algorithm 1 is the one we propose that obtains new samples  $\mathbf{x}^{new}$  drawn simultaneously from restaurants indexed by  $\mathbf{j}$  and associated seating arrangement  $\mathbf{s}^{new}$ , given previous samples  $\mathbf{x}^{old}$  and  $\mathbf{s}^{old}$ .

---

**Algorithm 1** MH-RCDSampler( $j, x^{old}, s^{old}$ ): Metropolis-Hastings sampler for restricted collapsed draw

---

```

1:  $s_L^{old} = s^{old}$ 
2: for  $\ell = L$  downto 1 do
3:    $s_{\ell-1}^{old} = \text{removeCustomer}(x_\ell^{old}, j_\ell, s_\ell^{old})$            { Remove customers for  $x_1^{old}, \dots, x_m^{old}$  sequentially from  $s^{old}$  }
4:    $r_\ell^{old} = p(x_\ell^{old}, s_{\ell-1}^{old})$                                { Calculate factors for accept ratio }
5: end for
6:  $s_0^* = s_0 = s_0^{old}$ 
7:  $x^* \sim q_x(x; s_0)$                                            { Draw  $x^*$  from proposal distribution  $q(x)$  of draws. }
8: for  $\ell = 1$  to  $L$  do
9:    $r_\ell^* = p(x_\ell^*, s_{\ell-1}^*)$                                    { Calculate factors for accept ratio }
10:   $s_\ell^* = \text{addCustomer}(x_\ell^*, j_\ell, s_{\ell-1}^*)$                  { Add customers for  $x_1^*, \dots, x_m^*$  sequentially to  $s_0^*$  }
11: end for
12:  $s^* = s_L^*$                                                    { Obtain proposal seating  $s^*$  }
13:  $R = \min\left(1, \frac{q_x(s^{old})}{q_x(s^*)} \prod_{\ell=1}^L \frac{r_\ell^*}{r_\ell^{old}}\right)$        { Calculate acceptance probability }
14: return  $\langle x^{new}, s^{new} \rangle = \begin{cases} \langle x^*, s^* \rangle & \text{with probability } R \\ \langle x^{old}, s^{old} \rangle & \text{otherwise.} \end{cases}$  { Accept/reject proposed sample }

```

---

The first half of this sampler is similar to a sampler for a single draw; it consists of removing old customers (line 3), choosing a new sample (line 7), and adding the customers again (line 10). The main difference is that there are  $L$  times of iteration for each call removeCustomer/addCustomer, and the calculation of  $r$ , which is later used for acceptance probability  $R$ .

## 4 Gibbs sampler for HCRP-HMM

This section describes a series of samplers for HCRP-HMM. First, we present the step-wise Gibbs sampler as the simplest example. After that, we describe a blocked Gibbs sampler using a forward-backward algorithm. We also explain the HCRP version of the beam sampler [?] as well as the split-merge sampler [?] for iHMM.

### 4.1 Step-wise Gibbs sampler

A step-wise Gibbs sampler for HCRP-HMM is easily constructed using an RCD sampler (Algorithm 5 in the Appendix describes one Gibbs sweep). We slightly modified the proposal distribution  $q(x_t)$  from that suggested in Section 3.2, in order to ensure that  $x_{t+1}$  is proposed with non-zero probability even when no table in  $s_0$  serves dish  $x_{t+1}$ :

$$q_x(x_t) \propto \left( p(x_t | s_0^{(x_{t-1})}) + \left( \frac{\alpha_0 \gamma}{(\alpha_0 + n_{x_{t-1} \cdot})(\gamma + m_{\cdot})} \right) \delta_{x_{t+1}} \right) \cdot p(x_{t+1} | s_0^{(x_t)}) \cdot p(y_t | F_0^{(x_t)}) \quad (33)$$

### 4.2 Blocked Gibbs sampler

We can construct an alternate sampling scheme under the framework of RCD sampler that resamples a block

of hidden states simultaneously, based on the forward-backward sampler [?]. The idea is that we run the forward-backward sampler with a predictive transition distribution from HCRP-HMM, and use the result as a proposal of restricted collapsed draw.

For iHMM, the forward-backward sampling algorithm [?] cannot be directly used, because the forward probability values for an infinite number of states have to be stored for each time step  $t$  [?]. This is not the case for HCRP-HMM, because predictive transition probability  $\hat{\pi}$  from given seating assignment  $s_0$ , which is given as Eqs. (31) and (32), only contains transition probability for finite number  $K$  of states plus one for  $k^{new}$ . Thus we only need to store  $K + 1$  forward probability for each time step  $t$ .

Result  $\bar{x}$  of the forward-backward sampler, however, cannot be used directly as the proposal; the  $i$ -th state of the proposal  $x_i^*$  is equal to  $\bar{x}_i$  when  $\bar{x}_i \neq k^{new}$ , but we need to assign new state indices to  $x_i^*$  whenever  $\bar{x}_i = k^{new}$ . In particular, when  $k^{new}$  has appeared  $W \geq 2$  times, all appearances of  $k^{new}$  may refer either to the same new state, or to  $W$  different states, or to anything in between the two, in which some appearances of  $k^{new}$  share a new state.

To achieve this purpose, we prepare special CRP  $Q^*$  that accounts for the previously unseen states, marked by  $k^{new}$  in the result of the forward-backward sampler. Specifically, each table in  $Q^*$  has a dish with an unused state index, and each appearance of  $k^{new}$  is replaced with a draw from  $Q^*$ . This construction ensures that every state sequence is proposed with a non-zero probability, and allows the proposal probability to be easily calculated. The concentration parameter of  $Q^*$  is set as equal to  $\gamma$ . To handle the case where some of the new states are equal to  $x_{t_{b+1}}$ , i.e., index of the state that succeeds to the resampling block, we add to  $Q^*$



an extra customer that corresponds to  $x_{t_{b+1}}$  when  $x_{t_{b+1}}$  does not appear in  $s_0$ ,

Resulting proposal probability is:

$$q_{\mathbf{x}}(\mathbf{x}^*) = \left( \prod_{\ell=0}^L \hat{\pi}_{\bar{x}_{\ell+1} \bar{x}_{\ell}} \cdot \prod_{\ell=1}^L F_{\bar{x}_{\ell}}(y_{\ell}) \right) \cdot \prod_{\ell: \bar{x}_{\ell} = k^{new}} p(x_{\ell}^* | Q^*) , \quad (34)$$

where the first factor accounts for the forward probability of the sequence, and the second factor accounts for probability of the new state assignment.

Note also that, to make a sensible proposal distribution, we cannot resample the whole state sequence simultaneously. We need to divide the state sequence into several blocks, and resample each block given the other blocks. The size of a block affects efficiency, because blocks that are too large have lower accept probability, while with blocks that are too small, the algorithm has little advantage over step-wise Gibbs sampling.

Algorithm 8 in the Appendix describes one sweep of a blocked Gibbs sampler for an HCRP-HMM.

### 4.3 Beam sampling

Beam sampling for HDP-HMM [?] is a sampling algorithm that uses slice sampling [?] for transition probability to extract a finite subset from the state space. Although the possible states are already finite in HCRP-HMM, the same technique may benefit sampling of HCRP-HMM by improving efficiency from the reduced number of states considered during one sampling step.

We just need replace the call to ForwardBackwardSampling in Algorithm 8 with the call to BeamSampling to use beam sampling with HCRP-HMM. A brief overview of the beam sampling is:

1. Sample auxiliary variables  $\mathbf{u} = (u_0, \dots, u_L)$  as  $u_{\ell} \sim \text{Uniform}(0, \pi_{x_{\ell} x_{\ell-1}})$ ,
2. For  $\ell = 1, \dots, L$ , calculate forward probability  $q(x'_{\ell} = k')$  using a slice of transition probability  $q(x'_{\ell} = k') = F_{k'}(y_{\ell}) \sum_k \mathbb{I}(\pi_{k'k} > u_{\ell-1}) q(x'_{\ell-1} = k)$ ,
3. For  $\ell = L, \dots, 1$ , sample the states  $x'_{\ell}$  backwardly, i.e.  $p(x'_{\ell} = k) \propto \mathbb{I}(\pi_{x'_{\ell+1}k} > u_{\ell})$ .

For details, refer to the original paper [?].

Some remarks may be needed for the calculation of  $q_{\mathbf{x}}^*$ , i.e., the proposal probability for the state sequence. Although beam sampling has a different proposal distribution from forward-backward sampling, we can use the same calculation of proposal probability used in acceptance probability as that of forward-backward sampling. This is be-

cause beam sampling satisfies the detailed balance equation, which ensures that the ratio of proposal probability with beam sampling  $\frac{q_{slice}^*}{q_{slice}}$  is always equal to the ratio of the probability obtained by forward-backward sampling  $\frac{q^*}{q_{fd}}$ .

### 4.4 Split-Merge Sampling

We can integrate the split-merge sampling algorithm, which is another sampling approach to Dirichlet process mixture models [?], into HCRP-HMM using the RCD sampler. A split-merge sampler makes a proposal move that tries to merge two mixture components into one, or to split a mixture component into two; the sampler then uses a Metropolis-Hastings step to stochastically accept the proposal. Based on the RCD framework, we can extend the split-merge sampler into HCRP, which can be applied to HCRP-HMM. Within the context of HMM, the sampler corresponds to merge two state indices into one, or to split a state index into two.

Our implementation is based on an improved version of the split-merge sampler, called the sequentially-allocated merge-split sampler [?], which produces a split proposal while sequentially allocating components in random order. To deal with temporal dependency in HMM, we identify fragments of state sequences to be resampled within the state sequence, and perform blocked Gibbs sampling for each fragment in random order.

We added one important optimization to the split-merge sampling algorithm. Since a merge move is proposed much more frequently than a split move, and the move has a relatively low accept probability, it is beneficial if we have a way of determining whether a merge move is rejected or not earlier. Let us point out that, when proposal probability for a merge move is calculated, the accept probability is monotonically decreasing. Consequently we sample  $R^{thr}$ , the threshold of accept probability, at the beginning of the algorithm and stop further calculation when  $R$  becomes less than  $R^{thr}$ . Algorithm 9 in the Appendix is the split-merge sampling algorithm for HCRP-HMM.

Split-merge sampling allows faster mixing when it is interleaved with other sampling strategies. We examine split-merge sampling with each of the samplers we have presented in this paper.

## 5 Experiments and Discussion

This section presents two series of experiments, the first with small artificial sequences and the second with a sequence of natural language words.

## 5.1 Settings

We put gamma prior  $\text{Gamma}(1, 1)$  on  $\alpha_0$  and  $\gamma$ , and sampled between every sweep using an auxiliary variable method [?] in all the experiments. We introduced HCRP as a prior of emission distributions as well, and its hyperparameters were also sampled in the same way.

The initial state sequence given to the sampler is the result of a particle filter with 100 particles.

We measured autocorrelation time (ACT) to evaluate mixing. Given a sequence of values  $\mathbf{x} = x_1, x_2, \dots, x_T$ , its mean  $\mu$  and variance  $\sigma^2$ ,  $ACT(\mathbf{x})$  are defined as follows:

$$ACF_t(\mathbf{x}) = \frac{1}{(T-t)\sigma^2} \sum_{i=1}^{T-t} (x_i - \mu)(x_{i+t} - \mu) \quad (35)$$

$$ACT(\mathbf{x}) = \frac{1}{2} + \sum_{t=1}^{\infty} ACF_t(\mathbf{x}) \quad . \quad (36)$$

Since with larger  $t$ ,  $ACF_t(\mathbf{x})$  is expected to converge to zero, we used  $ACF_t(\mathbf{x})$  for  $t \leq 1000$ .

For artificial sequence, we evaluated mutual information between the  $h_t$ , hidden state used in sequence generation and  $x_t$ , inferred states as follows:

$$MI = \sum_h \sum_x p(x, h) \log \frac{p(x, h)}{p(x)p(h)} \quad . \quad (37)$$

For natural language text, the inferred model is evaluated by multiple runs of a particle filter on a given test sequence of length  $T_{test}$ . We specifically construct a particle filter with  $Z = 100$  particles for each sampled model state  $s_z$ , and evaluate likelihood  $l(y_i | s_z)$  for each emission. Finally, we calculate the perplexity (the reciprocal geometric mean of the emission probabilities) of the test sequence:

$$PPL = \exp \left( -\frac{1}{T_{test}} \sum \log \hat{l}(y_i) \right) \quad (38)$$

where

$$\hat{l}(y_i) = \frac{1}{Z} \sum_{z=1}^Z l(y_i | s_z) \quad . \quad (39)$$

The samplers we chose for comparison are the step-wise Gibbs sampler with direct assignment representation [?], which uses stick-breaking for the root DP and CRP for the other DPs, the step-wise Gibbs sampler with stick-breaking construction, and the beam sampler with stick-breaking construction [?]. For fair comparison between different algorithms, we collected samples to evaluate the autocorrelation time and perplexity on a CPU-time basis (excluding the time used in evaluation). All the algorithms were implemented with C++ and tested on machines with an Intel Xeon E5450 at 3 GHz.




Figure 3: Automaton that generates Sequence 2. Circles denote hidden states, and the same alphabet emissions are observed from states within an oval group. A dashed arrow denotes transition with probability 0.8, a bold arrow denotes transition with probability 0.84, and a solid arrow denotes emission with probability 1/3.

## 5.2 Artificial data

The first series of experiments are performed with two small artificial sequences. Sequence 1 consists of repeating sequence of symbols A-B-C-D-B-C-D-E... for length  $T = 500$ , and we run the sampler 30 s for burn-in, and after that, a model state is sampled every 2 s until a total of 300 s is reached. Sequence 2 is generated from the simple finite state automaton in Figure 3 for length  $T = 2500$ , and we use 60 s for burn-in and total 600 s. We evaluated the mutual information between the inferred hidden states and the true hidden states.

Figure 4 shows the distribution of mutual information for 100 trials after 300 s. We can see that some of the samplers based on the proposed method achieved a better mutual information compared to existing samplers. The improvement depends on the type of sequence and the samplers.

For Sequence 1, we can see that split-merge sampling yields better results compared to other samplers. Although HMM with eight hidden states can completely predict the sequence, the samplers tend to be trapped in a local optimum with five states in the initial phase, because our selected prior of  $\gamma$  poses a larger probability on a smaller number of hidden states, Detailed investigations (Figure 5) confirmed this analysis.

For Sequence 2, on the other hand, blocked samplers worked very efficiently. Step-wise samplers generally worked poorly on the sequence, because the strong dependency on temporally adjacent states impedes mixing. Still, step-wise Gibbs sampler for HCRP-HMM outperformed the beam sampler with the stick-breaking process. The blocked Gibbs sampler had inferior performance due to its heavy computation for a large number of states, but the beam sampler for HCRP-HMM was efficient and performed well. Combination with a small number of split-merge samplers increases the performance (more split-merge sampling leads to lower performance by occupying

computational resource for the beam sampler). From averages statistics of samplers (Table 1), we can see that (1) the increase of mutual information cannot be described only by the increase of the number of states; (2) The accept ratio for the Gibbs trial has a very high accept rate; (3) Split-merge samplers have a very low accept rate, but still make improvement for mutual information.

### 5.3 Natural language text

We also tested the samplers using a sequence of natural language words from *Alice’s Adventure in Wonderland*. We converted the text to lower case, removed punctuation, and placed a special word EOS after every sentence to obtain a corpus with 28,120 words; we kept the last 1,000 words for test corpus and learned on a sequence with length  $T = 27120$ . We introduce a special word UNK (unknown) to replace every word that occurred only once, resulting in  $|\Sigma| = 1,487$  unique words in the text. We took 10,000 s for burn-in, and sampled a model state for every 120 s, until the total of 172,800 s. Table 2 summarize the averaged statistics for 18 trials.

We found that step-wise sampling outperformed blocked sampling (including beam sampling). The reason for this may be the nature of the sequence, which has a lower temporal dependency. Blocked Gibbs sampling, in particular, consumes too much time for one sweep to be of any practical use. We also found that split-merge sampling had a very low accept rate and thus made little contribution to the result.

Yet, we can see the advantage of using HCRP representation over stick-breaking representation. The direct assignment (DA) algorithm showed a competitively good perplexity, reflecting the fact that DA uses stick-breaking for only the root DP and uses the CRP representation for the other DP. Though step-wise Gibbs sampling and its slice sampling version seems outperforming DA slightly, we need to collect more data to show that the difference is significant. At least, however, we can say that now many sampling algorithms are available for inference, and we can choose a suitable one depending on the nature of the sequence.

## 6 Conclusion and Future Work


We have proposed a method of sampling directly from constrained distributions of simultaneous draws from a hierarchical Chinese restaurant process (HCRP). We pointed out that, to obtain a correct sample distribution, the seating arrangements (partitioning) must be correctly sampled for restricted collapsed draw, and we thus proposed applying the Metropolis-Hastings algorithm to the seating arrangements. Our algorithm, called the Restricted Collapsed Draw (RCD) sampler, uses a naïve sampler to provide a

proposal distribution for seating arrangements. Based on the sampler, we developed various sampling algorithms for HDP-HMM based on HCRP representation, including blocked Gibbs sampling, beam sampling, and split-merge sampling.


The applications of the RCD sampler, which is at the heart of our algorithms, are not limited to HCRP-HMM. The experimental results revealed that some of the proposed algorithms outperform existing sampling methods, indicating that the benefits of using a collapsed representation exceed the cost of rejecting proposals.

The main contribution of this study is that it opens a way of developing more complex Bayesian models based on CRPs. Since the RCD sampler is simple, flexible, and independent of the particular structure of a hierarchy, it can be applied to any combination or hierarchical structure of CRPs. Our future work includes using this algorithm to construct new Bayesian models based on hierarchical CRPs, which are hard to implement using a non-collapsed representation. Planned work includes extending HDP-IOHMM [?] with a three-level hierarchical DP (e.g., the second level could correspond to actions, and the third level, to input symbols).






(a) Sequence 1




(b) Sequence 2


Figure 4: Average mutual information of sampled hidden states




(a) HDP-HMM (SB) SGibbs




(b) HDP-HMM (DA) SGibbs




(c) HDP-HMM (SB) Beam



(d) HCRP-HMM Beam



(e) HCRP-HMM Beam  
Split-Merge 3/sweep



(f) HCRP-HMM Beam  
Split-Merge 25/sweep

Figure 5: Distribution of mutual information for Sequence 1. X-axis shows mutual information and Y-axis shows frequency. Block size  $\approx 6$  for HCRP-HMM Beam sampling.

Table 1: Experimental results for Sequence 2

name	MI	ACT	#states	#states	secs/sweep	Gibbs accept rate	SM accept rate
HDP-HMM (DA) SGibbs	2.92		0.527	14.910	0.044	—	—
HDP-HMM (SB) Beam	3.04		0.640	14.720	0.032	—	—
HCRP-HMM SGibbs	3.18		0.719	16.210	0.026	0.999666	—
HCRP-HMM SGibbs +SM=2	3.28		0.615	17.190	0.030	0.999632	0.000631
HCRP-HMM SGibbs +SM=13	3.19		0.493	16.830	0.044	0.999619	0.000650
HCRP-HMM SSlice	2.82		0.820	15.950	0.009	0.999847	—
HCRP-HMM SSlice +SM=2	2.86		0.705	17.330	0.013	0.999822	0.000827
HCRP-HMM SSlice +SM=13	2.59		0.604	16.400	0.030	0.999830	0.000910
HCRP-HMM BGibbs	3.01		0.317	14.900	0.206	0.995525	—
HCRP-HMM BGibbs +SM=2	3.12		0.513	16.270	0.237	0.995135	0.000985
HCRP-HMM BGibbs +SM=13	3.18		0.637	16.530	0.265	0.994715	0.000715
HCRP-HMM Beam	3.21		0.866	15.180	0.016	0.997233	—
HCRP-HMM Beam +SM=2	<b>3.37</b>		0.898	16.910	0.019	0.996369	0.000497
HCRP-HMM Beam +SM=13	3.28		0.875	17.070	0.034	0.996316	0.000532

DA: Direct Assignment SB: Stick-Breaking construction  
MI: Mutual Information ACT: Auto-correlation time, samples collected for every 0.1 s  
#states: number of states SGibbs: step-wise Gibbs  
SSlice: step-wise Gibbs with slice sampling (beam sampling with block size=1)  
BGibbs: blocked Gibbs (block size  $\approx 8$ )  
Beam: beam sampling (block size  $\approx 8$  for HCRP-HMM,  $T$  for stick-breaking)  
SM: Split-Merge sampler (+SM= $n$  denotes SM trials per Gibbs sweep)

Table 2: Experiments on Natural language text

Sampler	Perplexity	# states	sec/sweep	Gibbs accept rate	SM accept rate
HDP-HMM (DA) SGibbs	134.22	313.056	10.017	—	—
HDP-HMM (SB) SGibbs	151.10	242.389	38.045	—	—
HDP-HMM (SB) Beam	178.59	68.444	16.126	—	—
HCRP-HMM SGibbs	133.31	379.833	7.027	0.999861	—
HCRP-HMM SGibbs+SM=130	131.66	386.833	7.664	0.999857	0.000052
HCRP-HMM SGibbs+SM=5400	135.94	336.278	31.751	0.999880	0.000050
HCRP-HMM SSlice	131.17	422.000	0.469	0.999986	—
HCRP-HMM SSlice+SM=130	131.67	409.833	0.976	0.999993	0.000052
HCRP-HMM SSlice+SM=5400	152.76	254.722	36.617	0.999993	0.000056
HCRP-HMM BGibbs	199.14	1840.833	29380.261	0.992748	—
HCRP-HMM Beam	141.77	603.333	80.681	0.995627	—
HCRP-HMM Beam+SM=130	142.69	567.278	72.217	0.995612	0.000124
HCRP-HMM Beam+SM=5400	141.07	495.667	84.925	0.995554	0.000101

For HCRP-HMM, the block size  $\approx 10$ .

## A Miscellaneous Algorithms

---

**Algorithm 2** getProb( $j, k, s$ ): Calculate  $p(x_{ji} = k | s)$

---

```

if  $m_{\cdot k} = 0$  then
  return  $\frac{\alpha_0 \frac{\gamma}{m_{\cdot k} + \gamma}}{n_{j\cdot} + \alpha_0}$ 
else
  return  $\frac{n_{j\cdot k} + \alpha_0 \frac{m_{\cdot k}}{m_{\cdot k} + \gamma}}{n_{j\cdot} + \alpha_0}$ 
end if

```

---



---

**Algorithm 3** addCustomer( $j, k, s^{old}$ ): Adds new customer eating dish  $k$  to restaurant  $j$ .

---

```

 $s := s^{old}$ 
With probabilities proportional to:  $n_{jtk}$  ( $t = 1, \dots, m_{j\cdot}$ ): Increment  $n_{jtk}$  (the customer sits at  $t$ -th table)
 $\alpha_0 \frac{m_{\cdot k}}{m_{\cdot} + \gamma}$ : sit customer at a new table  $t^{new}$  serving dish  $k$  in restaurant  $j$  ( $n_{jt^{new}k} := 1, k_{jt^{new}} := k$ , increment  $m_{jk}$ )
return updated  $s$ 

```

---



---

**Algorithm 4** removeCustomer( $j, k, s^{old}$ ): Removes existing customer eating dish  $k$  from restaurant  $j$ .

---

```

 $s := s^{old}$ 
Sample  $t_{ji}$  in proportional to  $n_{jt_{ji}k}$ 
Decrement  $n_{jt_{ji}k}$  (the customer at  $t_{ji}$ -th table is removed)
if  $n_{jt_{ji}k}$  becomes zero then
  Remove the unoccupied table  $t_{ji}$  from restaurant  $j$ , decrement  $m_{jk}$ 
end if
return updated  $s$ 

```

---

## B Step-wise Gibbs sampler

To manipulate emission probability  $F(x_i)$  with a conjugate prior, we introduced a similar notation to HCRP, which can be intuitively understood.

---

**Algorithm 5** Step-wise Gibbs sweep for HCRP-HMM

---

**Input:**  $y_1, \dots, y_i$ : observed emissions  
 $x_1, \dots, x_i$ : previously inferred states  
 $s^{old}$ : set of CRP seating arrangements  
 $F^{old}$ : set of emission distributions

**for**  $i = 1, \dots, T$ , in random order **do**  
 $s_1 = \text{removeCustomer}(x_{i+1}, x_t, s^{old})$   
 $r_3^{old} = \text{getProb}(x_{i+1}, x_t, s_1)$   
 $F_0 = \text{removeCustomer}(y_i, x_i, F^{old})$   
 $r_2^{old} = \text{getProb}(y_i, x_t, s_1)$   
 $s_0 = \text{removeCustomer}(x_i, x_{i-1}, s_1)$   
 $r_1^{old} = \text{getProb}(x_i, x_{i-1}, s_0)$   
Sample  $x_i^*$  in proportion to  $q(x_t)$  where  

$$q(x_t) \propto \left( \frac{\alpha_0 \gamma}{(\alpha_0 + n_{x_{t-1} \dots})(\gamma + m \dots)} \delta_{x_{t+1}} + p(x_t | S_0^{(x_{t-1})}) \right) \cdot p(y_t | F_{x_t}) \cdot p(x_{t+1} | S_0^{(x_t)})$$
  
 $r_1^* = \text{getProb}(x_t^*, x_{t-1}, s_0)$   
 $s_1 = \text{addCustomer}(x_t^*, x_{t-1}, s_0)$   
 $r_2^* = \text{getProb}(y_t, x_t, F_0)$   
 $F^* = \text{addCustomer}(y_t, x_t, F_0)$   
 $r_3^* = \text{getProb}(x_{t+1}, x_t^*, s_1)$   
 $s^* = \text{addCustomer}(x_{t+1}, x_t^*, s_1)$   
 $R := \min \left( 1, \frac{r_1^*}{r_1^{old}} \frac{r_2^*}{r_2^{old}} \frac{r_3^*}{r_3^{old}} \frac{q(x_t)}{q(x_t^*)} \right)$   
 $\langle x_t, s, F \rangle := \begin{cases} \langle x_t^*, s^*, F^* \rangle & \text{with probability } R \\ \langle x_t, s^{old}, F^{old} \rangle & \text{otherwise} \end{cases}$

**end for**

---

## C Blocked Gibbs sampler

For details on the ForwardBackwardSampling routine, please refer to the literature [?].

---

**Algorithm 6**  $\text{removeSeq}(i_0, i_1, x, s^{old}, F^{old})$ : remove customers for a part of state sequence  $(x_{i_0}, \dots, x_{i_1})$

---

$L = i_1 - i_0 - 1$   
 $s_L = \text{removeCustomer}(x_{i_0+L}, x_{i_1}, s^{old})$   
 $r_L^{old} = p(x_{j_i} = k | s)$   
 $F_L = F^{old}$

**for**  $\ell = L - 1$  **downto** 0 **do**  
 $s_\ell = \text{removeCustomer}(x_{i_b+\ell}, x_{i_b+\ell-1}, s_{\ell+1})$   
 $F_\ell = \text{removeCustomer}(y_{i_b+\ell}, x_{i_b+\ell}, F_{\ell+1})$   
 $r_\ell^{old} = \text{getProb}(x_{i_b+\ell}, x_{i_b+\ell-1}, s_\ell)$   
 $\text{getProb}(y_{i_b+\ell}, x_{i_b+\ell}, F_\ell)$

**end for**  
**return**  $\langle s_0, F_0, \prod_{\ell=0}^L r_\ell^{old} \rangle$

---



---

**Algorithm 7**  $\text{addSeq}(i_0, i_1, x, s_0, F_0)$ : add customers for a part of state sequence  $(x_{i_0}, \dots, x_{i_1})$

---

$L = i_1 - i_0 - 1$   
**for**  $\ell = 0$  **to**  $L - 1$  **do**  
 $r_\ell^* = \text{getProb}(x_{i_b+\ell}, x_{i_b+\ell-1}, s_\ell)$   
 $\text{getProb}(y_{i_b+\ell}, x_{i_b+\ell}, F_\ell)$   
 $s_{\ell+1} = \text{addCustomer}(x_{i_b+\ell}^*, x_{i_b+\ell-1}^*, s_\ell)$   
 $F_{\ell+1} = \text{addCustomer}(y_{i_b+\ell}, x_{i_b+\ell}^*, F_\ell)$

**end for**  
 $r_L^* = \text{getProb}(x_{i_1}, x_{i_1-1}^*, s_L^*)$   
 $s^* = \text{addCustomer}(x_{i_1+L}, x_{i_1-1}^*, s_L^*); F^* = F_L^*$

**return**  $\langle s^*, F_L^*, \prod_{\ell=0}^L r_\ell^* \rangle$

---



---

**Algorithm 8** Blocked Gibbs sweep for HCRP-HMM

---

**Input:**  $y_1, \dots, y_i$ : observed emissions  
 $x = x_1, \dots, x_T$ : previously inferred states  
 $s$ : set of CRP seating arrangements  
 $F$ : set of emission distributions  
 $B$ : number of blocks

Choose block boundaries  $i_1, \dots, i_{B-1} \in \{2, \dots, T\}$ ;  
 $i_0 := 1, i_B = T$

**for**  $b = 0, \dots, B - 1$ , in random order **do**  
 $\langle s_0, F_0, r^{old} \rangle = \text{removeSeq}(i_b, i_{b+1} - i_b, x, s, F, 0)$

$x_i^* = x_i$  for all  $t < i_b$  or  $t \geq i_{b+1}$   
 $(x_{i_b}^*, \dots, x_{i_b+L-1}^*) = \text{FBSampler}(\hat{\pi}|_{S_0, F_0, y_{i_b:i_b+L-1}, x_{i_b-1}, x_{i_b+L}})$   
Calculate  $q^{old} = q(x_{i_b}, \dots, x_{i_b+L-1})$  and  $q^* = q(x_{i_b}^*, \dots, x_{i_b+L-1}^*)$   
 $Q^{old} = \text{CRP}(\gamma, H)$   
 $Q^* = \text{CRP}(\gamma, H)$

**if**  $x_{t_{b+1}}$  refers to a new state in  $s_0$  **then**  
 $Q^* := \text{addCustomer}(x_{t_{b+1}}, Q^*)$   
 $Q^{old} := \text{addCustomer}(x_{t_{b+1}}, Q^{old})$

**end if**  
**for**  $t = t_b$  **to**  $t_{b+1} - 1$  **do**  
**if**  $x_i$  refers to a new state in  $s_0$  **then**  
 $q^{old} := q^{old} * \text{getProb}(x_i, Q^{old})$   
 $Q^{old} := \text{addCustomer}(x_i, Q^{old})$

**end if**  
**if**  $x_i^* = k^{new}$  **then**  
sample  $s \sim Q^*$ ;  $x_i^* := s$   
 $q^* := q^* * \text{getProb}(x_i^*, Q^*)$   
 $Q^* := \text{addCustomer}(x_i, Q^*)$

**end if**  
**end for**  
 $S_0^* = S_0; F_0^* = F_0$   
 $\langle s^*, F^*, r^* \rangle = \text{addSeq}(i_0, L, x, s_0, F_0)$   
 $R := \min \left( 1, \frac{q^{old}}{q^*} \cdot r^{old} \cdot r^* \right)$   
 $\langle x, s, F \rangle := \begin{cases} \langle x^*, s^*, F^* \rangle & \text{with probability } R \\ \langle x^{old}, s^{old}, F^{old} \rangle & \text{otherwise} \end{cases}$

**end for**

---

## D Split-Merge sampler

---

### Algorithm 9 Split-Merge Sampler for an HCRP-HMM

---

**Input:**  $y_1, \dots, y_T$ : observed emissions  
 $x_1, \dots, x_T$ : previously inferred states  
 $\mathbf{s}^{old}$ : set of CRP seating arrangements  
 $\mathbf{F}^{old}$ : set of emission distributions  
 $R^{thr} \sim \text{Uniform}(0, 1)$   
 Choose distinct  $t_1, t_2 \in \{1, \dots, T\}$   
 Identify all fragments  $(b_i, e_i)$  s.t. for all  $t \in (b_i, \dots, e_i)$ ,  $x_i \in \{x_{t_1}, x_{t_2}\} \wedge t \notin \{t_1, t_2\}$ , and not contained in other fragments  
 Permute fragments randomly  
 Let  $U$  be the number of fragments  
 $\mathbf{s}_{U+1} = \mathbf{s}^{old}, \mathbf{F}_{U+1} = \mathbf{F}^{old}$   
**if**  $x_{t_1} = x_{t_2}$  **then** { Try split move }  
     **for**  $i = U$  **downto** 1 **do**  
          $\langle \mathbf{s}_i, \mathbf{F}_i, r_i^{old} \rangle = \text{removeSeq}(b_i, e_i, \mathbf{x}, \mathbf{s}_{i+1}, \mathbf{F}_{i+1})$   
          $q_i^{old} = 1$   
     **end for**  
      $x_{t_2}^* = \text{new } k \text{ index}$   
**else** { Try merge move }  
     **for**  $i = U$  **downto** 1 **do**  
          $\langle \mathbf{s}_i, \mathbf{F}_i, r_i^{old} \rangle = \text{removeSeq}(x_{b_i:e_i}, \mathbf{s}_{i+1}, \mathbf{F}_{i+1})$   
          $q_i^{old} = \frac{\text{SeqProb}(\hat{\pi}_{|s_{i+1}}, \mathbf{F}_i, y_{b_i:e_i}, x_{b_i-1:e_i+1})}{\text{ForwardProb}(\hat{\pi}_{|s_{i+1}}, \mathbf{F}_i, y_{b_i:e_i}, x_{b_i-1}, x_{e_i+1}; \{x_{t_1}, x_{t_2}\})}$   
     **end for**  
      $x_{t_2}^* = x_{t_1}$   
**end if** { Remove customers that accounts for transitions around  $x_{t_2}^{old}$  }  
 $\mathbf{F}_0 = \text{removeCustomer}(y_{t_2}, x_{t_2}, \mathbf{F}_1)$   
 $p_0^{old} = p(y_{t_2} | \mathbf{F}_{x_{t_2}})$   
 $\mathbf{s}_0 := \mathbf{s}_1$   
**if**  $t_2 - 1$  is not in any fragment **then**  
      $\mathbf{s}_0 := \text{removeCustomer}(x_{t_2}, x_{t_2-1}, \mathbf{s}_0)$   
      $r_0^{old*} = \text{getProb}(x_{t_2}, x_{t_2-1}, \mathbf{s}_1)$   
**end if**  
**if**  $t_2 + 1$  is not in any fragment **then**  
      $\mathbf{s}_0 := \text{removeCustomer}(x_{t_2} + 1, x_{t_2}, \mathbf{s}_0)$   
      $r_0^{old*} = \text{getProb}(x_{t_2+1} x_{t_2}, \mathbf{s}_0)$   
**end if**  
 $q_0^{old} = q_0^* = 1$   
 (continue to Algorithm 10)

---



---

**Algorithm 10** Split-Merge Sampler for an HCRP-HMM (continued)

---

{ Add customers that accounts for transitions around  $x_{t_2}^*$  }

$$p_0^* = p(y_{t_2} | \mathbf{F}_{x_{t_2}^*}^*)$$

$$\mathbf{F}_1^* = \text{addCustomer}(y_{t_2}, x_{t_2}^*, \mathbf{F}_0^*)$$

$$\mathbf{s}_1^* := \mathbf{s}_0$$

**if**  $t_2 + 1$  is not in any fragment **then**

$$r_0^* = \text{getProb}(x_{t_2+1} | x_{t_2}, \mathbf{s}_1^*)$$

$$\mathbf{s}_1^* := \text{addCustomer}(x_{t_2} + 1, x_{t_2}, \mathbf{s}_1^*)$$

**end if**

**if**  $t_2 - 1$  is not in any fragment **then**

$$r_0^* = \text{getProb}(x_{t_1} | x_{t_2-1}, \mathbf{s}_1^*)$$

$$\mathbf{s}_1^* := \text{addCustomer}(x_{t_2}, x_{t_2} - 1, \mathbf{s}_1^*)$$

**end if**

**if**  $x_{t_1} = x_{t_2}$  **then**

{ Try split move }

**for**  $i = 1$  to  $U$  **do**

$$x_{b_i}^*, \dots, x_{b_i+L_i-1}^* = \text{LimitedFBSampler}(\hat{\pi}|_{s_{i+1}^*}, F_i^*, y_{b_i:b_i+L_i-1}, x_{b_i-1}^*, x_{b_i+L_i}^*; \{x_{t_1}^*, x_{t_2}^*\})$$

$$\text{SeqProb}(\hat{\pi}|_{s_{i+1}^*}, F_i^*, y_{b_i:e_i}, x_{b_i-1:e_i+1}^*)$$

$$q_i^* = \frac{\text{SeqProb}(\hat{\pi}|_{s_{i+1}^*}, F_i^*, y_{b_i:e_i}, x_{b_i-1:e_i+1}^*)}{\text{ForwardProb}(\hat{\pi}|_{s_{i+1}^*}, F_i^*, y_{b_i:e_i}, x_{b_i-1}^*, x_{e_i+1}^*; \{x_{t_1}^*, x_{t_2}^*\})}$$

$$\langle \mathbf{s}_{i+2}^*, \mathbf{F}_{i+1}^*, r_i^* \rangle = \text{addSeq}(b_i, e_i, \mathbf{x}^*, \mathbf{s}_{i+1}^*, \mathbf{F}_i^*)$$

**end for**

**else**

{ Try merge move }

**for**  $i = 1$  to  $U$  **do**

$$R^{cur} = \prod_{i'=0}^{i-1} \frac{r_{i'}^*}{q_{i'}^*} \cdot \prod_{i=0}^I \frac{r_i^{old}}{r_i^{old}}$$

**if**  $R^{thr} \geq R^{cur}$  **then**

rejection determined, exit loop

**end if**

$$x_{b_i}^*, \dots, x_{e_i}^* = x_{t_1}$$

$$q_i^* = 1$$

$$\langle \mathbf{s}_i^*, \mathbf{F}_{i-1}^*, r_i^* \rangle = \text{addSeq}(b_i, e_i, \mathbf{x}, \mathbf{s}_{i+1}, \mathbf{F}_i)$$

**end for**

**end if**

$$R = \prod_{i=0}^I \frac{r_i^*}{r_i^{old}} \cdot \prod_{i=1}^I \frac{q_i^{old}}{q_i^*}$$

$$\langle \mathbf{x}, \mathbf{s}, \mathbf{F} \rangle = \begin{cases} \langle \mathbf{x}^*, \mathbf{s}^*, \mathbf{F}^* \rangle & R^{thr} < R \\ \langle \mathbf{x}^{old}, \mathbf{s}^{old}, \mathbf{F}^{old} \rangle & \text{otherwise} \end{cases}$$


---